SINGLE-CHIP FPGA-BASED PROCESSOR-SENSOR-CONTROLLER

W. Schmidt[(1)] , A.Mälkki[(2)]

[(1)]Finnish Meteorological Institute, P.O.Box 503, 00101 Helsinki, Walter.Schmidt@fmi.fi
[(2)]Finnish Meteorological Institute, P.O.Box 503, 00101 Helsinki, Anssi.Malkki@fmi.fi

ABSTRACT

For planetary exploration the available resources mass, volume and energy are severely restricted while the environmental conditions like radiation, temperature range and pressure can be very demanding. At the same time the suitable components are less and less available. In this article we show several options to meet these requirements, implemented mainly on ESA's SMART-1 mission and on the Lander Philae of ESA´s comet mission ROSETTA. The strategy employed is to replace as many components as possible by building blocks inside radiation tolerant Field Programmable Gate Arrays. RISC-Processor, measurement sequencer, power optimization and conversions between analog and digital signals all co-resided on the same chip. The concept allows the development of single-chip instrument electronics.

1. INTRODUCTION

The space division of the Finnish Meteorological Institute in Helsinki has been involved in the design, manufacturing and operation of space-based instruments since the late seventies. The activities concentrated mainly on the control electronics including software for plasma instruments, the electrical ground support equipment including flight operation support and on the development of meteorological sensors for planetary atmospheres.
During the last decade the availability of processors as well as analog and digital components suitable for space exploration has decreased while the requirements such components have to meet became more demanding. Only in the field of Field Programmable Gate Arrays (FPGA) significant progress was made offering an ever increasing number of different radiation hardened or tolerant devices also for space applications. Inside ESA's space program this was utilized for the first time on a larger scale with the comet mission Rosetta. In the following missions like Mars Express or Smart-1 this technology became standard.
The requirements especially for planetary exploration are partly contradictory:
- Hardened against large accumulated radiation doses but low mass budget for shielding
- Operational at extreme low and/or high temperatures
- High demands on processing speed but low power consumption
- Recovery possibility from many failure causes without requiring extra resources.

In this article we present different low-cost approaches to these requirements as utilized in the Spacecraft Potential, Electron and Dust Emission Analyzer (SPEDE) [1] on board ESA´s SMART-1 satellite and in the Permittivity Probe (SESAME/PP) [2] on the Rosetta Lander PHILAE. In PP special power saving strategies were implemented inside the radiation hardened FPGA, while SPEDE, developed several years later, realized also most of the analog and data processing functionality inside one FPGA. For design validation, operation and data evaluation an integrated approach for the electrical ground support equipment of both instruments was chosen which implemented the complete FPGA functionality inside a laptop.
Currently already available space qualified FPGAs allow the implementation of all these building blocks and concepts inside the same chip, reducing mass, volume and power consumption to a minimum and allowing the implementation of radiation hazard avoidance techniques without excess shielding needs like majority-voting logic, error-correction codes etc.

2. FPGA-BASED LOW-POWER PROCESSOR

SMART-1 (see fig.1) was launched with a mass of 350kg and a payload of 15kg from Kourou on September 23, 2003 with the launcher Ariane-5. On Sunday 3.9.2006 at about 05:42 GMT, the probe crashed into a volcanic plain with an impact speed of about 7,200 km/h. The satellite's main purpose was the operational testing of an ion propulsion engine. Targeted to the Moon it mapped in detail the Moon surface with several instruments. During the cruise phase SPEDE was part of the plasma monitoring instrumentation observing the plasma environment around the spacecraft with and without activated engine.

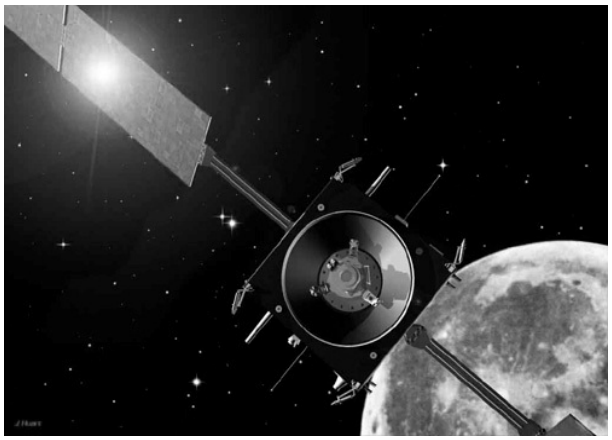

Figure1. The SMART-1 spacecraft at the Moon

The SPEDE command and data processing was handled by a 16-bit processor realized as an FPGA building block. The implementation followed to some extend the ideas presented by John Rible from Sand Piper Technologies / California in 1999[3] as his QS5 Baby RISC FPGA implementation of a CPU. Each instruction takes 2 clock cycles, during which in parallel the next two program bytes are read from the program memory, organized as 8bit x 512k EEPROM. The data RAM had the same organization. The processor had 8 different registers which needed 3 bit inside the 16-bit wide instruction for addressing. Register 0 was the directly testable status register with separate bits indicating 0 or negative results after each instruction. In our implementation the other bits were mapped to external events like sequencer and telecommunication status information, power monitor and EEPROM writing time-out, saving asynchronous interrupt handling and special hardware access. A block diagram of the complete electronics is shown in fig. 2.

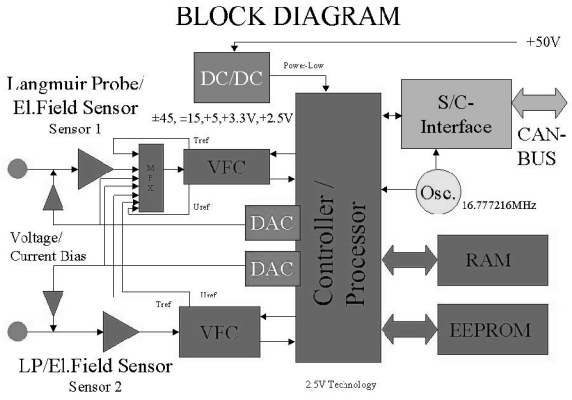

Figure 2. SPEDE block diagram

2.1. Program integrity

On power-up or commanded instrument reset the processor executed the boot instruction hardcoded inside the FPGA. This instruction was implemented as an absolute jump to an address page defined via a 3-bit counter, which gave access to 8 different program pages. Starting at 0 after switch-on, the program on page 0 was executed, starting with a checksum test of the complete code. If this failed, it ended in an endless loop, forcing the hardware watch-dog circuitry to reset the processor and increment the boot-page counter. The processor would now start on page two and so on.
If during program execution a single-event upset changed the program memory, the watch-dog reset would recover the instrument and continue with another program copy. With a special telecommand any of the program pages could be selected directly allowing different program versions to co-reside inside the instrument.
Any part of the memory or complete code pages could be replaced by a sequence of memory load commands, allowing the optimization of the code for different mission phases. Also the duplicating of one program page to another was possible with a single command to correct possible radiation-caused code errors without the need of a complete software upload.

2.2. Memory management

To improve the protection against accidental overwriting of program memory, program and data address space was separated by hardware. The 16-bit program address space accessible through the 16-bit registers was reduced to 14 bit, freeing the 2 most significant address bits for memory management control. The used radiation hardened EEPROM and the RAM were divided into pages of 16 Kbyte size each.
The processor provided an additional input-output port (I/O) with up to 64 different addresses, which could be accessed with a single instruction. This separate address space was used for data transfer with the communication interface, the measurement sequencer, the Real-Time-clock, watch-dog circuit, a special EEPROM overwrite protection hardware and the page registers.
The page registers formed the core of the memory management. 6-bit wide, they allowed access to any page available inside EEPROM or RAM. An additional bit extended the external memory to the boot program area, implemented inside the FPGA. The 8th available bit separated the program page register from the data page register. While a change to the data page register became directly available for the next instruction, the new program page value was buffered without changing the current program page.

The page change was initiated by one of the two free program address bits. When set, the program continued on the new page at the address defined in the instruction. At the same time the hardware saved the previous page back to the page buffer, allowing a simple cross-page return from a subroutine. This value could be read back by the program in case nested cross-page subroutine jumps were used.
An abbreviated "long-jump" instruction allowed the unconditional jump directly to another page. In this case 6 bit of the jump instruction were used as new program page register contents directly forced into the program page register without the need of a preceding I/O-instruction. As only 8 address bits remained, the program continued on an address with the lowest 6 bits set to 0. This feature was mainly used for the boot process and system program parts controlling the writing to the EEPROM, which had therefore to be executed in RAM.
Each processor instruction contained one bit initiating directly after instruction execution an unconditional or conditional return from a subroutine. Together with the page management this allowed the generation of an extremely compact code. The final flight software including telemetry formatting and on-line data analysis fit into one page of 16 Kbyte!

3. FPGA-BASED ADC/DAC

For SPEDE it was essential that measurements could be carried out at times without the solar array generating power while at the same time being inside the Earth's radiation belt. These requirements excluded most of the available converters between analog and digital signals, as radiation hardened analog circuits tend to have high power consumption. Instead, the analog voltages were generated under program control via an external combined serial-parallel resistor network, followed by a low-power radiation hardened operational amplifier (fig.3). Forcing different FPGA-pin combinations to logical '0', a linear 8-bit DAC-function was simulated. Any precision, dynamic range and characteristics can be implemented by choosing the appropriate number of control bits, resistor values and suitable control logic. The conversion time depends only on the used controller frequency and may be as short as a few nanoseconds while the needed power is only a fraction of a μW.

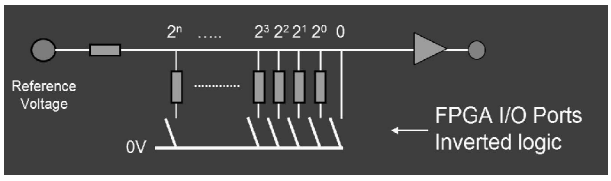

Figure 3 Resistor-network based DAC (simplified)

Positive and negative analog signals were measured via commercial radiation hardened Voltage–to-Frequency converters (VFC) for up to 200 kHz. The available circuits convert input voltages into symmetrical square wave signals. The linear dependency between voltage and frequency is temperature stabilized., the power consumption is in the order of a few mW. The resulting pulses were then either counted in an adjustable time window defining the used integration time, or the length of a set of pulses was compared with the processor clock. The latter method provides better resolution for small signals around the center frequency, while the former gives a large dynamic range. A combination of pulse counting and pulse length measurement was used in FMI's pressure probe on Cassini/Huygens' HASI instrument (fig. 4) giving a large dynamic range for the measured voltage and at the same time a good resolution for small signals.

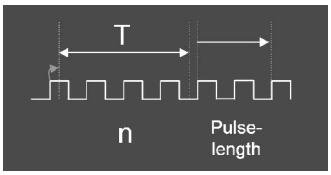

Figure 4. Frequency measurement based ADC

4. POWER SAVING STRATEGIES

In the SPEDE instrument, the most effective power saving means were the replacement of ADC, DAC and processor by implementations inside the same FPGA. The instrument itself had a maximum power consumption of 200mW with only radiation hardened pA analog amplifiers, one RAM and one EEPROM chip besides the radiation tolerant FPGA. The complete electronics with the power supply and communication boards underneath the main board is shown in fig. 5.

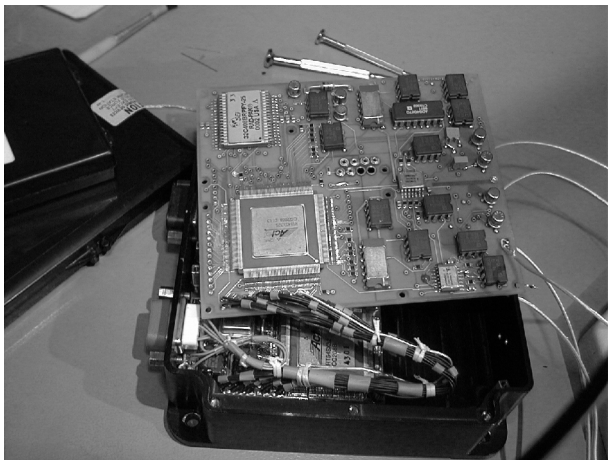

Figure 5. SPEDE Electronics board

For the PP-instrument on Rosetta's Lander Philae a different strategy is used to adapt the power needs to the different phases of operation. This is crucial for the operation during comet night without the availability of solar energy. The electronics is divided into four functional groups, differently shaded in fig. 6.
The power to the three different analog parts (transmitter with DAC and current monitors, differential receiver with multiplexer for housekeeping signals and an integrator as plasma wave detector) can be activated independently from one another. The main power consumption is in the transmitter part, which is kept off most of the time.

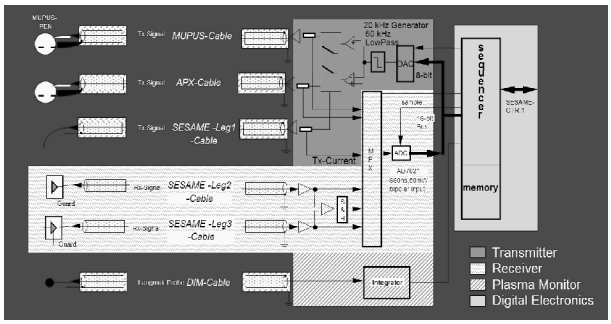

Figure 6. Rosetta/PP Electronics block diagram

The digital controller part consists of an FPGA and two memory chips in CMOS technology. Their power demands can be significantly reduced by reducing the clock speed or operating the FPGA as completely static logic. By using separate clock networks which can be separately controlled via some FPGA logic functions, the power consumption can be optimized to the different operational phases. In stand-by mode the logic statically decodes incoming commands while the clock input is disabled. The needed power is then only about 50 µW. In this mode all FPGA-internal control registers can be loaded and read. If for memory access a clock signal is needed, this is enabled for that network participating in the memory control activities. During measurements other parts of the sequencer logic become active until the predefined measurement cycle has finished, when the clock network to the respective network will be automatically disabled again. Fig. 7 demonstrates the principle of this logic implementation.

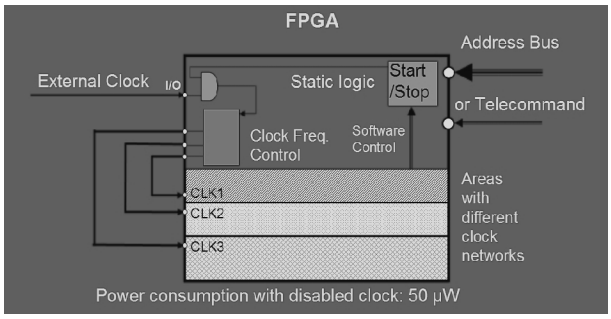

Figure 7. Clock control as means for power reduction

The external clock signal is not directly connected to any clock network but to a standard input of the FPGA. Static logic can distribute the clock signal, reduce it to different frequencies at the same time or disable the clock altogether. The various pre-conditioned clock signals are routed via standard output buffers externally back into the FPGA's dedicated clock input ports, driving the different available clock networks.

5. OPERATION AT EXTREME TEMPERATURE

Only a few commercially available components can operate outside the military temperature range of -55C to +125C. In many cases it is not the electronics itself, incorporated inside the die of the circuit, which limits the useful temperature range, but the carrier towards low temperature and the used bonding technique towards high temperatures and large mechanical shocks. A possible solution to these problems is to integrate as many functions as possible into the same die to reduce the number of failure prone inter-component connections. If as carrier material a ceramics optimized for the intended temperature range and special welding techniques for the bonding process are used, the temperature range can be extended to below -100°C and above 200°C with standard silicon-based dies. This approach was successfully used by the team which developed the APX spectrometer on board Mars Pathfinder and the Rosetta Lander [4].

6. INSTRUMENT CONTROL & SIMULATION

A highly integrated system needs good verification tools. For both presented FPGA-based controllers the complete FPGA code was converted into C-language computer code, implemented on a LINUX-operating-system based computer. An external interface allowed connecting a representative model of the instrument with a standard USB-port on the same computer. Program switches allowed selecting either the simulator code or directly the hardware to carry out user commands. A TKL-based graphical user interface exchanged ASCII-coded information with this C-Program via a network-port based in-

terface. This modular approach allowed the physical separation of the operator console from the actual hardware as long as an internet connection could be established.

All aspects from FPGA-code verification and flight software optimization to flight operation control and data analysis could be tested with the same environment. As the instrument was from the beginning treated as an integrated system, the different aspects could be easily moved between FPGA implementation, flight software functionality and operation and data analysis tasks, until an optimal balance was found. The Rosetta/PP simulator code part was extended to allow directly the simulation of different electrical comet surface properties. The system then performed the "measurements" and generated telemetry packets, which were used to verify and optimize the data analysis software of the ground segment.

Fig. 7 shows the user interface for the SPEDE instrument during a calibration measurement. All the information in the bottom half of the display corresponds to command and configuration parameters, while the upper half shows the measurement data of both sensors in semi-logarithmical scale. Other display options are possible.

A direct-command interface allows to put either the flight software or even the FPGA-code execution into single-step debugging mode or to set breakpoints, inspect memory areas, activate automatic timing commands or patch the running software with a command identical to the one used for the instrument in space.

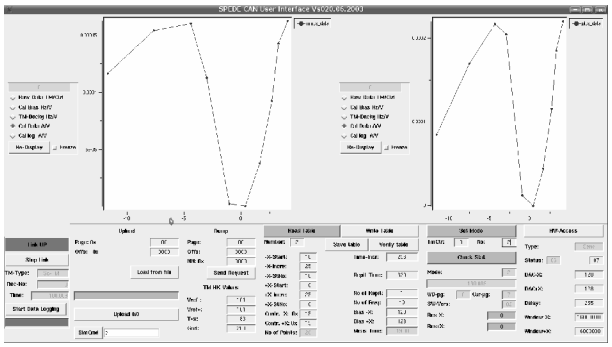

Figure 7. SPEDE EGSE and simulator

Like any high-level commercial software debugger the simulator gives directly the link between just executed binary command code and the related software source code line. A simulated clock linked to the actual execution time of each command allows absolute timing of measurements and generates correct time stamps inside the telemetry packets.

The Rosetta/PP signal simulator allowed to add different noise and offset levels to the simulated physical measurement, separately for each of the two channels (see fig. 8). This was used to optimize the stability of the onboard analysis algorithm which extracts directly the physical properties from the measured data series thereby reducing the needed telemetry by a factor of 2000.

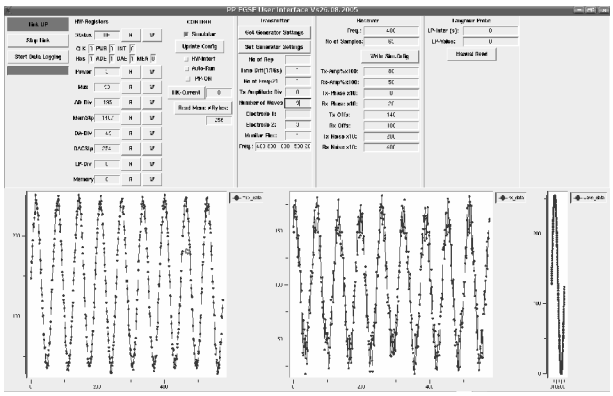

Figure 8. PP signal simulation including noise

7. TOWARDS A SINGLE-CHIP CONTROLLER

Possible applications: In applications which are mass, volume and power critical or are intended for difficult environmental conditions like planetary Landers, balloons etc, this approach looks promising. Mass, volume and power demands can be further reduced by combining the remaining analog functions with the central controller in a single-chip hybrid or ASIC solution. The presented approach is especially suitable for instrument providers with limited budget and manufacturing resources, which might make the purchase and qualification of custom-made radiation hardened components prohibitive. The available resources can in stead be concentrated on the development and optimization of the detectors.

8. REFERENCES

1. Mälkki, A., Schmidt, W., Laakso, H., Grard, R., Escoubet, C.P., Wahlund, J.-E., Blomberg, L., Marklund, G. and Johlander, B., The SPEDE experiment on SMART-1: Instrument, mission, and science objectives, *Geophysical Research Abstracts*, Vol. 5., 10004, 2003

2. K. J. Seidensticker, D. Möhlmann, I. Apathy, W. Schmidt et al., Sesame, An Experiment of the Rosetta Lander Philae: Objectives and General Design, *Space Science Reviews*, Vol. 128, 301-337, 2007

3. J.Rible, Guided Exploration of two FPGA-based CPU-Designs, http://www.sandpipers.com/cpuclass.html, 1999

4. G. Klingelhöfer, The Rosetta Alpha Particle X-Ray Spectrometer (APXS), *Space Science Reviews*, Vol. 128, no 1-4, 2007